Fill it up: Exploiting partial dependency annotations in a minimum spanning tree parser.

Liang Sun[1] **Jason Mielens**[2] **Jason Baldridge**[2]

[1]Department of Mechanical Engineering
The University of Texas at Austin
sally722@utexas.edu

[2]Department of Linguistics
The University of Texas at Austin
{jmielens,jbaldrid}@utexas.edu

arXiv:1611.08765v1 [cs.CL] 26 Nov 2016

Abstract

Unsupervised models of dependency parsing typically require large amounts of clean, unlabeled data plus gold-standard part-of-speech tags. Adding indirect supervision (e.g. language universals and rules) can help, but we show that obtaining small amounts of direct supervision—here, partial dependency annotations—provides a strong balance between zero and full supervision. We adapt the unsupervised ConvexMST dependency parser to learn from partial dependencies expressed in the Graph Fragment Language. With less than 24 hours of total annotation, we obtain 7% and 17% absolute improvement in unlabeled dependency scores for English and Spanish, respectively, compared to the same parser using only universal grammar constraints.

1 Introduction

Unsupervised parsing solutions are simultaneously an attractive yet troublesome method for handling low-data scenarios. The performance of unsupervised parsers has increased dramatically in recent years (Klein and Manning, 2004; Naseem et al., 2010), making them a potentially viable option for constructing labeled corpora on limited budgets. However, their performance is often outmatched by small amounts of labeled data (Blunsom and Cohn, 2010; Spitkovsky et al., 2012). Further, recent work using linguistically-informed error analysis on unsupervised Combinatory Categorial Grammar parsing shows that entire syntactic phenomena are outside the scope of existing unsupervised parsers (Bisk and Hockenmaier, 2015). Accordingly, most recent work in this area has focused on methods of providing sources of indirect annotation, whether via linguistic world-knowledge (Naseem et al., 2010; Grave and Elhadad, 2015), partial annotations (Flannery et al., 2011; Mielens et al., 2015) or cross-lingual information transfer (Naseem et al., 2012).

With unsupervised parsing, data collection is not entirely eliminated: a large amount of clean, relevant data is needed. Also, evaluations of unsupervised techniques typically rely on gold part-of-speech tags. Obtaining clean data for many languages is actually a difficult process–complicated by issues such as language identification, digitization, and varying or absent orthographies. This challenge also exists in many domain adaptation scenarios.

We explore the effectiveness of creating small amounts of labeled data using the Graph Fragment Language (GFL), an annotation scheme designed for speed and ease (Schneider et al., 2013; Mordowanec et al., 2014). We create 270 English and 2297 Spanish partial sentence annotations using GFL, using a mix of expert and non-expert annotators. We then adapt the minimum spanning tree based parsing technique of Grave & Elhadad (2015) to use these partial annotations in addition to universal dependency rules it already exploits. Throughout this work we will refer to this parser as ConvexMST.[1]

We present parsing results with and without gold part-of-speech tags. When using predicted POS tags, our experiments show that exploiting cheap, incomplete direct supervision in addition to language universals provides large absolute performance im-

[1]Code available at github.com/jmielens/convex-mst

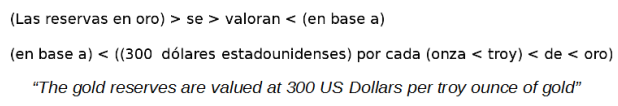
(Las reservas en oro) > se > valoran < (en base a)
(en base a) < ((300 dólares estadounidenses) por cada (onza < troy) < de < oro)
"The gold reserves are valued at 300 US Dollars per troy ounce of gold"

Figure 1: Spanish GFL Example: Parentheses indicate a constituent-style bracket, angle brackets indicate direct dependency relations.

provements for both English and Spanish: 6.3% for the former and 17.3% for the latter. Furthermore, the ConvexMST parser dramatically outperforms the Gibbs sampler parser of Mielens et al. (2015) using the supervision (English: +5.2%; Spanish: +14.4%). We also show that the extra supervision provided by gold POS tags heavily influences results; in particular, it inflates performance when only using language universals. Experiments that rely on gold POS tags alone are thus not reliable indicators of performance in true low-resource settings.

2 Annotation

2.1 Graph Fragment Language

We use the Graph Fragment Language (GFL) (Schneider et al., 2013) to allow for light-weight, simple annotations that our annotators can easily learn and use confidently. The choice of annotation scheme is particularly important: we seek to optimize annotation speed rather than full-specification or high accuracy. In previous studies, the use of GFL has allowed for annotation rates of 2-3 times that of traditional dependency annotations while still maintaining a useful level of annotation density (Mordowanec et al., 2014; Mielens et al., 2015). Hwa (1999) demonstrated that it is most effective to provide high-level sentence constituents to a parser and allow it to fill in the low-level information itself.

The GFL annotation in Figure 1 shows two distinct notations. Constituent brackets are specified by parentheses and direct dependencies by angle brackets. Many words and phrases are underspecified. Allowing partial annotations dramatically increases the speed at which annotators can work, while simultaneously reducing error rates. These two effects both arise from being able to leave difficult or tedious portions of a sentence unspecified.

2.2 Filling in Partial Dependencies

A partial annotation produces a set of dependency tree fragments. Compared to an unlabeled sentence, this can substantially reduce the work a parser must do. When working with partial dependencies, there are two paths that can be taken with regard to overall model-building. In a 'Fill-then-Parse' setup, the partial dependencies are first filled-in to produce full dependencies that are then used to train a standard dependency parser. In a 'Fill+Parse' setup, one model both fills in and parses new sentences.

We use a Fill+Parse setup, while previous work focused on Fill-then-Parse. The major benefit of the former is that learning can be sensitive to the source of an arc in the training data—e.g., whether it came from an annotator or a universal rule. Fill-then-Parse obscures this distinction and not knowing how trustworthy an arc is can lead to additional errors. Indeed, Fill+Parse method produces better results for our datasets than Fill-then-Parse (see Section 4.2).

2.3 Simulated Cost Comparison

Many factors influence the cost of creating a corpus. Our goal is to minimize cost relative to the performance of a parser trained with the corpus. The actual cost of finding and paying annotators is the most obvious factor, and it will typically be higher for a low-resource language or highly specialized domain. Using a light-weight partial annotation scheme like GFL has the potential to increase the pool of qualified annotators and alleviate this challenge.

Given a partial annotation scheme like GFL, an additional cost factor is that of obtaining a particular level of completion for each sentence. Consider that for any sentence there are both 'low-hanging fruit' dependencies such as determiner attachment, and more difficult dependencies such as preposition attachment and long-distance relations. Harder dependencies take longer to annotate (and thus cost more), so it is worth considering cost metrics that incorporate completion percentage. In the absence of timing/expense data, we can simulate this intuition with a variable cost model for which each an additional dependency annotated in a sentence is more expensive than the previous one.

Figure 2 demonstrates the impact of completion cost. Parsing accuracies (for our parser introduced in

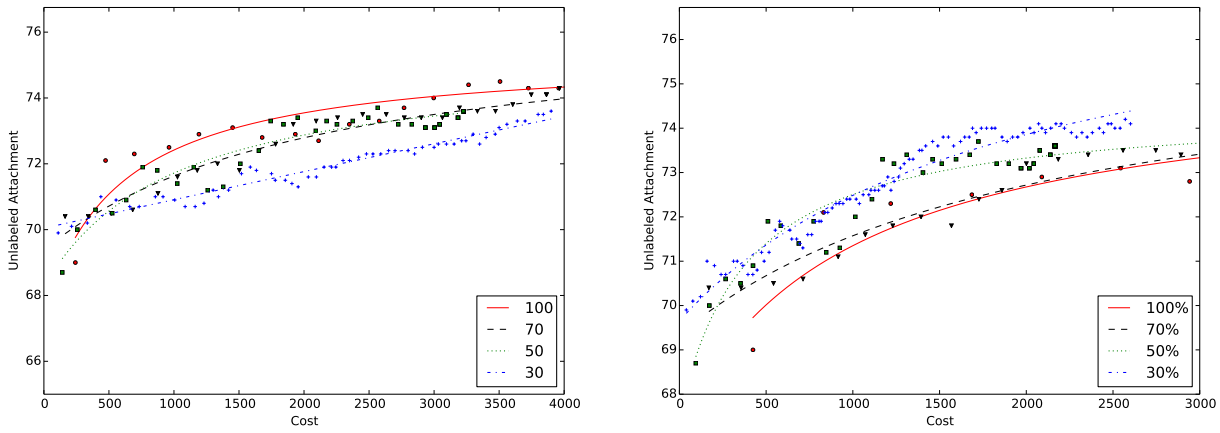

(a) Equal Cost

(b) Variable Cost

Figure 2: Comparison of performance versus total cost, Equal Cost is the sum of all specified dependencies, Variable Cost weights dependencies by completion percentage. Run on Spanish data using simulated partial dependencies.

the next section) are shown at different costs, under (a) simple equal (per arc) cost and (b) variable cost. We simulated the construction of various corpora by deriving partial dependencies from gold standard annotations), and show the cost curves for different sentence completion rates. 100% completion produces the best performance with equal costs, but under the more realistic variable cost model, 30% and 50% completion win. We show later that this pattern holds under actual timed annotation.

Garrette (2015) demonstrated the benefit of partial annotations for CCG parsing. They focused on the number of (partial) bracket annotations (as a proxy for annotation time), holding this fixed while varying the number of sentences. Strikingly, they found that having 40% of brackets across the full dataset was better than full brackets for 80% of the corpus. This result uses an equal cost-per-bracket assumption, so the difference would be even more favorable to partial annotations with a variable cost.

2.4 Unsupervised vs. Partial Annotations

Without any direct annotations, we must rely on indirect supervision such as universal grammar rules, cross-lingual information transfer, and domain adaptation. Following Grave & Elhadad (2015), we use the universal grammar rules in Table 1. Indirect supervision via these rules is achieved by biasing produced trees to conform to the rules. This is the only form of dependency supervision considered by Grave & Elhadad, though they do provide additional direct supervision via gold part-of-speech tags.

Verb $\mapsto$ Verb	Noun $\mapsto$ Noun
Verb $\mapsto$ Noun	Noun $\mapsto$ Adj
Verb $\mapsto$ Pron	Noun $\mapsto$ Det
Verb $\mapsto$ Adv	Noun $\mapsto$ Num
Verb $\mapsto$ Adp	Noun $\mapsto$ Conj
Adj $\mapsto$ Adv	Adp $\mapsto$ Noun

Table 1: Universal Grammar Rules

2.5 Data

We use two sources of data. To compare with prior work, we use the universal treebanks (version 2.0), which cover ten languages from a variety of language families (McDonald et al., 2013). We obtained GFL annotations for a subset of the English data, originally from WSJ Section 03 of the Penn Treebank, and we use simulation techniques to produce partial dependencies for the other languages.

Our second data source is the Spanish dependency treebank from the AnCora corpus (Taulé et al., 2008). For 1410 unique sentences of AnCora, we have partial dependencies specified in GFL by twelve annotators. Most sentences received a single partial annotation from a single annotator, but one section of the corpus was annotated by all annotators. As the original corpus is fully-specified for

gold dependencies, we can measure annotator agreement with a gold standard.

The background and experience of the annotators varied considerably. Roughly one third were native Spanish speakers, with the rest ranging from fluent non-native speakers to a few with just a single year of formal study. This was done intentionally to provide a large variance in the types and quality of annotations that they were able to provide.

Each annotator was trained for just 30 minutes. The nature of the annotations was explained and a small number of guidelines were provided. For instance, annotators were told that typically adjectives are dependents of nouns, nouns are dependents of verbs, and so on. These guidelines amounted to a summary of the rules in Table 1. During the annotation sessions, annotators were told to ask as many clarifying questions as needed, although in practice they needed very little guidance. Post-experiment debriefing interviews suggested that the straightforward nature of the GFL notation was very helpful and became clear within a few example sentences.

Despite minimal training time, annotators were able to produce relatively consistent annotations that agreed in large part with other annotators. Table 3 shows both pair-wise and overall agreement between annotators when considering arcs that each of the annotators in the pair had provided a head for. Overall agreement was high, with most pairwise numbers in the 70-80's, and agreement for individual annotators to the group is even higher – mostly in the 80's.

The partial annotation task proved helpful in terms of speed; our annotators were able to cover 750 tokens/hr, which compares favorably to the processes of the Penn Treebank, which achieved rates of 750-1000 tokens/hr for English (Marcus et al., 1993), and 300-400 tokens/hr for Chinese (Xue et al., 2005), both making use of initial parse suggestions from an existing parser. Efforts not using an existing parser proceed even slower; for instance the Ancient Greek Dependency Treebank reported rates of 100-200 tokens/hr (Bamman and Crane, 2011).

2.6 POS-Tagging

Our goal is to minimize real-world costs associated with producing a finished parsing model. To this end, we trained our own POS taggers using type label annotations (Garrette and Baldridge, 2013) rather than using gold-standard tags. We use universal POS tags rather than the finer-grained sets the source corpora use, both for simplicity and cross-language comparisons (Petrov et al., 2011).

	Partial EN	Full ES	Partial ES
Unique Sentences	270	135	1410
Total Sentences	270	135	2162
Number of Annotators	2	1	12
Total Annotation Hours	8	13	72

Table 2: Training Set Statistics

We trained taggers for all languages using a limited amount of the available gold data—ensuring that the accuracy is comparable with low-resource human-sourced taggers. We extract types from the corpus, rank them by frequency, and take the most frequent types to train the tagger. The cutoff on how many types to take is derived from the number of types the annotators in Garrette et al. (2013) were able to produce in two hours. The taggers all obtain around 80% accuracy.

3 Method

3.1 Convex-MST

This section provides a brief overview of the core parsing algorithm; for full details, see Grave & Elhadad (2015). We begin by considering a binary vector $\mathbf{y}$ that encodes all of the dependencies in our corpus, such that $\mathbf{y}_{ijk} = 1$ if sentence i has an arc with dependent j and head k. This representation leads to the problem formulation in Equation 1, where Y is the convex hull of all the valid tree assignments for $\mathbf{y}$, $\mathbf{n}$ is the number of possible dependency arcs in the corpus, $\mathbf{u}$ is a penalty vector that penalizes potential dependency arcs that are not in the set of universal dependency rules, and $\mathbf{w}$ is a weight vector learned during training

$$\min_{\mathbf{y} \in Y} \min_{\mathbf{w}} \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2 + \frac{\lambda}{2} \|\mathbf{w}\|_2^2 - \mu \mathbf{u}^T \mathbf{y} \quad (1)$$

This problem can be solved using Algorithm 1 (Grave and Elhadad, 2015).

3.2 Partial Dependency Features

The main modification we make is to add an additional term to penalize arcs that disagree with partial

Annotator	1	2	3	4	5	6	7	8	9	10	11	12	Avg.
1	1	0.73	0.9	0.88	0.55	0.77	1	0.94	0.9	0.28	0.95	0.67	.80
2	0.73	1	0.78	0.83	0.95	0.62	0.77	0.75	1	0.27	0.8	0.85	.78
3	0.9	0.78	1	0.85	0.64	0.8	0.9	0.82	0.96	0.3	0.85	0.72	.79
4	0.88	0.83	0.85	1	0.6	0.88	0.83	0.92	1	0.33	0.91	0.68	.81
5	0.55	0.95	0.64	0.6	1	0.46	0.64	0.55	0.83	0.23	0.59	0.85	.66
6	0.77	0.62	0.8	0.88	0.46	1	0.88	0.74	0.88	0.16	0.75	0.6	.71
7	1	0.77	0.9	0.83	0.64	0.88	1	0.94	1	0.2	1	0.67	.82
8	0.94	0.75	0.82	0.92	0.55	0.74	0.94	1	1	0.36	0.94	0.7	.81
9	0.9	1	0.96	1	0.83	0.88	1	1	1	0	1	0.81	.87
10	0.28	0.27	0.3	0.33	0.23	0.16	0.2	0.36	0	1	0.11	0.12	.28
11	0.95	0.8	0.85	0.91	0.59	0.75	1	0.94	1	0.11	1	0.68	.80
12	0.67	0.85	0.72	0.68	0.85	0.6	0.67	0.7	0.81	0.12	0.68	1	.70
Total	.85	.86	.86	.9	.82	.77	.98	.92	.96	.61	.94	.82	

Table 3: Pair-wise and total agreement by annotator. The 'Total' row shows agreement with the set of all other annotators and the 'Avg.' column is the average pairwise agreement.

Algorithm 1 Optimization algorithm from Grave & Elhadad (2015)

for $r \neq 0$ **do**
  Compute the optimal $\mathbf{w}$:
  $\mathbf{w}_t = \arg\min_{\mathbf{w}} \frac{1}{2n}\|\mathbf{y}_t - \mathbf{X}\mathbf{w}\|_2^2 + \frac{\lambda}{2}\|\mathbf{w}\|_2^2$
  Compute the gradient w.r.t. $\mathbf{y}$:
  $\mathbf{g}_t = \frac{1}{n}(\mathbf{y}_t - \mathbf{X}\mathbf{w}_t) - \mu\mathbf{u}$
  Solve the linear program:
  $\mathbf{s}_t = \min_{\mathbf{s}\in Y} \mathbf{s}^T\mathbf{g}_t$
  Take the Franke-Wolfe step:
  $\mathbf{y}_t = \gamma_t\mathbf{s}_t + (1-\gamma_t)\mathbf{y}_t$

annotations. Let $\mathcal{S}$ be the set of all indices on $\mathbf{y}$ where that head-dependent pair conforms to one of the universal rules. Then we can require that some proportion of the arcs in the corpus satisfy:

$$\frac{1}{n}\sum_{i\in\mathcal{S}} y_i \geq c$$

This is equivalent to $\mathbf{u}^T\mathbf{y} \geq c$, where:

$$u_i = \begin{cases} 1/n, & \text{if } i \in \mathcal{S}. \\ 0, & \text{otherwise.} \end{cases}$$

This is how the penalty term $\mu\mathbf{u}^T\mathbf{y}$ from (1) is derived. Similarly, we can add another penalty term that ensures a certain percentage of the arcs conform to the arcs specified by annotators. If we let $\mathcal{G}$ be the set of all indicies on $\mathbf{y}$ where the word pair conforms to the GFL annotations, then it is simple to construct an additional penalty term $\xi\mathbf{v}^T\mathbf{y}$.

There is a slight difference between the GFL penalty term and the universal rule penalty term. Whereas the universal rule penalty is based simply on whether the arc conforms or does not conform to the rules, the GFL annotations naturally lead to a three-way distinction: the annotation can specify that an arc *should* be present, *should not* be present, or make no commitment.

Accordingly, we modify $\mathcal{G}$ to be two sets, $\mathcal{G}_w$ and $\mathcal{G}_b$, where $\mathcal{G}_w$ is the set of all indicies on $\mathbf{y}$ where the word pair should have an arc, and $\mathcal{G}_b$ is the set of all indicies on $\mathbf{y}$ where the word pair should *not* have an arc. We refer to these as the whitelist and blacklist, accordingly. Under this formulation, the GFL-based penalty term $\xi\mathbf{v}^T\mathbf{y}$ is now made with:

$$v_i = \begin{cases} 1/n, & \text{if } i \in \mathcal{G}_w \\ -1/n, & \text{if } i \in \mathcal{G}_b \\ 0, & \text{otherwise} \end{cases}$$

This leads to the modified objective function in (2), which now seeks to find a solution that minimizes the number of arcs that violate both universal rules and the annotator-specified fragments.

$$\min_{\mathbf{y}\in Y}\min_{\mathbf{w}} \frac{1}{2n}\|\mathbf{y}-\mathbf{X}\mathbf{w}\|_2^2 + \frac{\lambda}{2}\|\mathbf{w}\|_2^2 - \mu\mathbf{u}^T\mathbf{y} - \xi\mathbf{v}^T\mathbf{y} \tag{2}$$

When no GFL annotations are specified for the corpus, the GFL penalty term goes to zero and the objective function reverts to its original formulation.

Specific arcs are added to $\mathcal{G}_w$ and $\mathcal{G}_b$ in a number of ways, based on the different types of GFL annotation. Consider the GFL annotation in Figure 3. Here, the annotator has specified a direct dependency with 'passed' as the head of 'congress'. The

congress has not passed < (a comprehensive plan)
passed < congress

Figure 3: GFL Whitelisting vs. Blacklisting

arc 'passed ← congress' is added to $\mathcal{G}_w$, while all other arcs of the form 'X ← congress' are added to $\mathcal{G}_b$ because 'congress' may only have a single head.

Brackets may also result in additions to the whitelist and blacklist. In Figure 3, 'a comprehensive plan' is bracketed. In this case, no arcs can be whitelisted, but many can be blacklisted. For instance, no word external to the bracket may be headed by a word in the bracket. This means arcs such as 'plan ← congress' must be in $\mathcal{G}_b$.

Also, 'passed' is indicated as the head of the entire bracket. We cannot whitelist any specific arcs with this information (since we do not know the head of the bracketed expression), but we know that no word internal to the bracket is headed by any word external to it, other than 'passed'. Hence, arcs such as 'congress ← plan' must be in $\mathcal{G}_b$.

4 Experiments and Discussion

We consider both simulated and actual partial annotations. Results based on actual annotation are the most important as they provide our best measure of performance under a realistic annotation setting. However, our Spanish annotators had only six hours each, and there was no inter-annotator communication or creation of annotation conventions, and no attempt to have them adopt the conventions in the gold-standard AnCora dependencies we evaluate against. Because of this, we include simulation results to eliminate this source of divergence to better measure the effectiveness of different methods for filling in missing arcs in a partial annotation. It of course also allows us to measure this for all the languages in the Universal Dependencies treebanks.

We consider three different supervision settings for ConvexMST:

- *UG* uses just the universal grammar based features, which is equivalent to the method used by Grave & Elhadad (2015).
- *GFL* uses just the human specified features.
- *GFL+UG* uses both.

	Degradation		
Language	None	Light	Heavy
DE	69.6	69.5	68.2
EN	79.5	77.5	72.8
ES	78.0	76.5	71.5
FR	82.6	82.3	75.7
IT	82.0	81.8	77.4
PT-BR	80.6	80.0	72.4
SV	77.7	77.1	76.9
Average	78.5	77.8	73.6

Table 4: Simulated degradation results. *Light* has 40% of arcs removed and *Heavy* has 70% removed.

Parser	Features	Gold Tags		Predicted Tags	
		EN	ES	EN	ES
RB	N/A	17.1	28.0	17.1	28.0
Gibbs	GFL	60.2	65.3	55.8	52.7
ConvexMST	UG	63.1	63.5	56.9	50.0
	GFL	65.9	70.5	61.2	67.1
	UG+GFL	**68.2**	**71.3**	**63.2**	**67.3**

Table 5: Directed dependency accuracy on English and Spanish universal treebanks using annotator provided GFL annotations, 10 or fewer words.

These three methods correspond with $\xi \neq 0$, $\mu \neq 0$, and $\xi\mu \neq 0$ in Equation 2. The training sets correspond with the 'Partial EN' and 'Partial ES' sets from Table 2. The set of sentences annotated with GFL is used as the training set for the *GFL*, *UG*, and *GFL+UG* methods.

4.1 Simulated partial dependencies

Simulated partial dependencies are produced by removing dependencies via a stochastic process that approximates how we instructed human annotators to focus their efforts. Arcs are removed top-down, with arcs lower in the tree being more likely to be deleted. This results in trees with more high-level structures and less lower-level information. Figure 4 demonstrates the stability of our parser under varying levels of such gold tree degradation. Missing arcs were recovered using our parse imputation scheme (using GFL+UG features), and the resulting parser was applied to the evaluation sentences. Accuracy decreases slightly to around 60% removal, and then degrades more rapidly after that. Table 4 provides numeric data for the simulations.

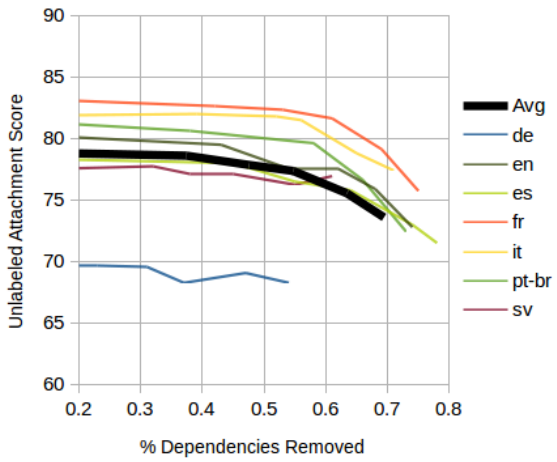

Figure 4: Degradation Simulations

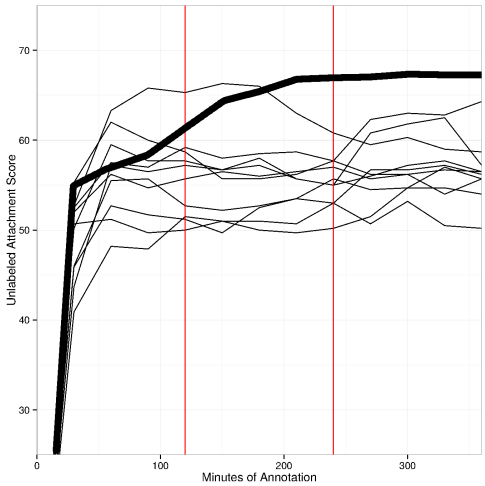

Figure 5: Learning curves for individual annotators (thin lines) and conglomerated training sets (thick line) over annotation time. Vertical lines indicate annotation session breaks.

4.2 Annotator-sourced partial dependencies

Table 5 gives semi-supervised parsing results on the English and Spanish treebanks for sentences with 10 or fewer words. To investigate the impact of POS taggers on parsing results, we conducted two series of experiments using POS tags trained by our own tagger as discussed in Section 2.6 (*Predicted Tag*) and gold POS tags extracted from treebank (*Gold Tag*). We compare against a right-branching baseline and the Gibbs parser of Mielens et al. (2015).

All the parsing methods handily beat the right-branching baseline. ConvexMST-UG (the model of Grave and Elhadad (2015)) beats the Gibbs parser with gold POS tags, but the ranking switches with predicted POS tags. This shows the effectiveness of ConvexMST, but highlights its brittleness with respect to tagging errors: bad tags lead to poor guidance from language universals. ConvexMST-GFL easily beats both these approaches: it exploits partial annotations much more effectively than the Gibbs parser and learns effectively without language universals. The difference is especially marked for predicted POS tags: ConvexMST-GFL beats ConvexMST-UG by 4.3% for English and 17.1% for Spanish. (Recall that there were 8 hours of annotation for English and 72 hours for Spanish.)

The best method of all uses both partial annotations and language universals: ConvexMST-UG+GFL improves on ConvexMST-GFL for both languages and POS conditions. The impact of the combination is greater for English, which has less GFL annotation. Overall, these results show that this combination is robust to varying amounts of partial annotations: the UG constraints are strong on their own and provide a strong basis without annotations, they contribute when there are not many annotations available, and eventually become less essential (but remain unharmful) as more are provided.

It is important to recall that the GFL annotations have no specific conformity to the gold standards of either *original* corpus. Our goal was to understand the overall behavior of different methods given the same free-wheeling, diverse annotations; it is likely that higher numbers would have been achieved had we guided annotators to use corpus conventions, or used full annotations provided by our annotators as the evaluation set. The former defeats the spirit of our exercise, and we did not have sufficient budget for the latter.

For Spanish, we also considered the performance of individual annotators alongside the full training set. The learning curves for individual annotators are shown in Figure 5. There is substantial variation in the curves for the individual annotators; however, the curve based on the union of all annotations at

each time step is smooth and is better than any individual past the three hour mark. One way to consider this is in terms of building an accurate parser quickly with multiple, diverse annotators, where wall clock time matters. Another way is to consider robustness with respect to possibly bad annotators. The next obvious steps would be to use active learning and to detect disagreement in annotators to either drop some or intervene to improve their quality. (Again, keep in mind that we are considering a "cold start" to this process, so there can be no gold standard for checking annotator quality.)

Comparison to Full Annotation To this point, all performance comparisons have been between different parse feature sets; we have demonstrated that the GFL features are complimentary to the UG features, and that when standing alone the GFL features are stronger than the UG features. The question of whether it might be more effective to simply have annotators produce full annotations is not addressed by these comparisons. To answer this question, we had our most experienced annotator fully annotate the same section that the other annotators did partially. Producing these full annotations required roughly 13 hours of time from the single expert annotator. In comparison, the other annotators were able to partially annotate the same section in roughly two hours each – a total of 24 hours. However, the theoretical wall clock time of the group of annotators could be as low as two hours if the sessions were run in parallel. These different training sets were once again used to train ConvexMST models that were evaluated on a held out test set. Table 6 contains the results of this experiment, demonstrating that the group of inexperienced annotators producing partial annotations was able to achieve similar performance levels to the single annotator producing full annotations. It should be noted that this comparison does not weight the results using the extrinsic costs associated with the production of the training data. In a real-world environment, the expert annotator would likely be more expensive than the inexperienced annotators, and possibly all of them combined (especially in a crowd-sourcing scenario). This makes the performance per unit cost for partial annotators even higher than Table 6 indicates. See Section 2.3 for discussion and modeling of these extrinsic cost effects.

Feature Set	Partial Annotations	Full Annotations
UG	56.9	58.8
GFL	61.2	62.8
GFL+UG	**63.2**	**66.6**

Table 6: Comparison between full and partial annotations, 10 or fewer words, using predicted POS tags.

4.3 Longer Sentences

We also evaluated ConvexMST with longer sentences: those with 20 words or less. For this, the right-branching baseline is 25.8%. When using all the annotations on the common set for all annotators, the scores for ConvexMST with UG, GFL, and GFL+UG are 47.6%, 54.4%, and 55.3%, respectively. The values are worse than for shorter sentences, as expected, but the pattern observed in Table 5 still holds: GFL annotations best UG alone, and their combination is the best of all.

4.4 Discussion & Error Analysis

POS-Tagging Impact We thought it important to consider imperfect POS-taggings because this entire framework is based off of the assumption that the user is working from essentially no pre-existing resources. Assuming the availability of gold-standard POS tags is antithetical to this idea, and is one way in which direct supervision can show up in otherwise unsupervised (or indirectly supervised) systems.

Many tagger errors are not likely to cause major problems during parsing; for instance mislabeling pronouns as nouns, or adverbs as adjectives, is unlikely to lead to major structural issues. However, more unlikely errors can cause more dramatic effects, as shown in Figure 6. Here, the phrase 'beating politically' (gold tags 'NOUN ADV') is mistagged as 'ADJ VERB', leading to the attachment of 'politically' to the root word and the reorganization of a substantial chunk of the sentence.

Weighting Constraint Violations For feature sets with both GFL and UG-based constraints, a weighting factor can bias the parser towards being more likely to respect either GFL or UG constraints. We experimented with this, and found that for the

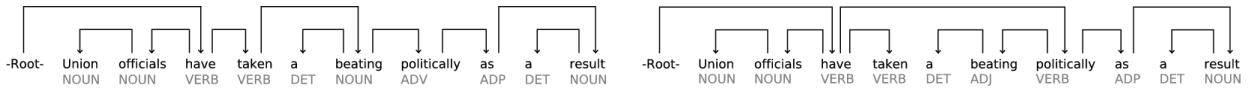
-Root- Union officials have taken a beating politically as a result
NOUN NOUN VERB VERB DET NOUN ADV ADP DET NOUN
-Root- Union officials have taken a beating politically as a result
NOUN NOUN VERB VERB DET ADJ VERB ADP DET NOUN

(a) Gold Tags (b) Predicted Tags

Figure 6: Differences in parsing results due to minimal POS tagging errors.

datasets we considered, the best results were obtained when we weighted violations of GFL constraints as worse than violations of UG constraints. This result is not entirely unexpected given the relative performances of the constraints on their own, but it provides more evidence that direct supervision even in small amounts can beat indirect supervision.

5 Conclusion

We have shown that human-sourced partial annotations can be exploited to learn effective dependency parsers in short period of time. The ConvexMST method we adapt from Grave and Elhadad easily combines constraints from both language universals and partial annotations, providing greater robustness from starting annotation until one runs out of budget or time. We demonstrate this with actual annotations produced for English and Spanish, using annotators with a range of experience.

Overall, we present a case for working in realistic settings by paying close attention to the various sources of annotation and tracking the real costs associated with that supervision. We believe that over-reliance on creeping supervision of this type may lead to an inaccurate picture of the cross-lingual and low-resource applicability of various models, and are encouraged by recent work on character-based models by Gillick et al. (2015) and Ballesteros et al (2015), among others. Their work shows viable models can be produced without relying on having annotations a priori, but rather learning representations on the fly that need not conform to any one set of standards.

Acknowledgments

Supported by the U.S. Army Research Office under grant number W911NF-10-1-0533. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the view of the U.S. Army Research Office.

References

[Ballesteros et al.2015] Miguel Ballesteros, Chris Dyer, and Noah A Smith. 2015. Improved transition-based parsing by modeling characters instead of words with LSTMs. *arXiv preprint arXiv:1508.00657*.

[Bamman and Crane2011] David Bamman and Gregory Crane. 2011. The ancient Greek and Latin dependency treebanks. In *Language Technology for Cultural Heritage*, pages 79–98. Springer.

[Bisk and Hockenmaier2015] Yonatan Bisk and Julia Hockenmaier. 2015. Probing the linguistic strengths and limitations of unsupervised grammar induction. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.

[Blunsom and Cohn2010] Phil Blunsom and Trevor Cohn. 2010. Unsupervised induction of tree substitution grammars for dependency parsing. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1204–1213. Association for Computational Linguistics.

[Flannery et al.2011] Daniel Flannery, Yusuke Miayo, Graham Neubig, and Shinsuke Mori. 2011. Training dependency parsers from partially annotated corpora. In *IJCNLP*, pages 776–784.

[Garrette and Baldridge2013] Dan Garrette and Jason Baldridge. 2013. Learning a part-of-speech tagger from two hours of annotation. In *HLT-NAACL*, pages 138–147. Citeseer.

[Garrette2015] Dan Garrette. 2015. *Inducing Grammars from Linguistic Universals and Realistic Amounts of Supervision*. Ph.D. thesis, University of Texas at Austin.

[Gillick et al.2015] Dan Gillick, Cliff Brunk, Oriol Vinyals, and Amarnag Subramanya. 2015. Multilingual language processing from bytes. *Arxiv preprint*.

[Grave and Elhadad2015] Edouard Grave and Noémie Elhadad. 2015. A convex and feature-rich discriminative approach to dependency grammar induction. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1375–1384, Beijing, China, July. Association for Computational Linguistics.

[Hwa1999] Rebecca Hwa. 1999. Supervised grammar induction using training data with limited constituent information. In *Proceedings of the 37th annual meeting of the Association for Computational Linguistics on Computational Linguistics*, pages 73–79. Association for Computational Linguistics.

[Klein and Manning2004] Dan Klein and Christopher D Manning. 2004. Corpus-based induction of syntactic structure: Models of dependency and constituency. In *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics*, page 478. Association for Computational Linguistics.

[Lei et al.2014] Tao Lei, Yu Xin, Yuan Zhang, Regina Barzilay, and Tommi Jaakkola. 2014. Low-rank tensors for scoring dependency structures. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 1381–1391.

[Marcus et al.1993] Mitchell P Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of English: The Penn Treebank. *Computational linguistics*, 19(2):313–330.

[McDonald et al.2013] Ryan T McDonald, Joakim Nivre, Yvonne Quirmbach-Brundage, Yoav Goldberg, Dipanjan Das, Kuzman Ganchev, Keith B Hall, Slav Petrov, Hao Zhang, Oscar Täckström, et al. 2013. Universal dependency annotation for multilingual parsing. In *ACL (2)*, pages 92–97. Citeseer.

[Mielens et al.2015] Jason Mielens, Liang Sun, and Jason Baldridge. 2015. Parse imputation for dependency annotations. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1385–1394, Beijing, China, July. Association for Computational Linguistics.

[Mordowanec et al.2014] Michael T. Mordowanec, Nathan Schneider, Chris Dyer, and Noah A. Smith. 2014. Simplified dependency annotations with GFL-Web. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 121–126, Baltimore, Maryland, USA, June. Association for Computational Linguistics.

[Naseem et al.2010] Tahira Naseem, Harr Chen, Regina Barzilay, and Mark Johnson. 2010. Using universal linguistic knowledge to guide grammar induction. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1234–1244. Association for Computational Linguistics.

[Naseem et al.2012] Tahira Naseem, Regina Barzilay, and Amir Globerson. 2012. Selective sharing for multilingual dependency parsing. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 629–637. Association for Computational Linguistics.

[Petrov et al.2011] Slav Petrov, Dipanjan Das, and Ryan McDonald. 2011. A universal part-of-speech tagset. *arXiv preprint arXiv:1104.2086*.

[Schneider et al.2013] Nathan Schneider, Brendan O'Connor, Naomi Saphra, David Bamman, Manaal Faruqui, Noah A Smith, Chris Dyer, and Jason Baldridge. 2013. A framework for (under) specifying dependency syntax without overloading annotators. *arXiv preprint arXiv:1306.2091*.

[Spitkovsky et al.2012] Valentin I Spitkovsky, Hiyan Alshawi, and Daniel Jurafsky. 2012. Three dependency-and-boundary models for grammar induction. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 688–698. Association for Computational Linguistics.

[Taulé et al.2008] Mariona Taulé, Maria Antònia Martí, and Marta Recasens. 2008. AnCora: Multilevel annotated corpora for catalan and spanish. In *LREC*.

[Xue et al.2005] Naiwen Xue, Fei Xia, Fu-Dong Chiou, and Marta Palmer. 2005. The Penn Chinese TreeBank: Phrase structure annotation of a large corpus. *Natural language engineering*, 11(02):207–238.